

Google Cloud Professional Cloud Network Engineer Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	16

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 - 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

- 1. Which feature distinguishes Traceroute from Ping?**
 - A. It measures packet loss**
 - B. It shows the average time for each hop**
 - C. It tests bandwidth limits**
 - D. It performs DNS lookups**
- 2. If a VM has 2 or fewer vCPUs, what is the maximum number of NICs it can have?**
 - A. 1**
 - B. 2**
 - C. 4**
 - D. 8**
- 3. What is the maximum response size that can be cached by Cloud CDN when byte-range requests are supported?**
 - A. 10 MB**
 - B. 5 TB**
 - C. 1 GB**
 - D. 50 MB**
- 4. What does the term 'interconnectivity options' in a VPC refer to?**
 - A. Methods for enhancing data storage**
 - B. Ways to connect the VPC with external networks**
 - C. Protocols for internal communications**
 - D. Settings for firewall configurations**
- 5. What does Google Cloud Armor provide to its users?**
 - A. Firewall protection against unauthorized access**
 - B. Security policies to protect applications from DDoS attacks**
 - C. Encryption for data at rest**
 - D. User authentication services**

6. Which command would you use to import a zone file into GCP?

- A. gcloud dns import zone-file**
- B. gcloud dns record-sets import**
- C. gcloud dns zone import**
- D. gcloud dns records add**

7. What is a limitation of the Netperf tool compared to IPerf?

- A. It cannot measure latency**
- B. It has less customizable features**
- C. It does not support UDP testing**
- D. It only works in Windows environments**

8. Which of the following best describes firewall rules?

- A. Protocols to manage application performance**
- B. Configurations that allow secure connections to public networks**
- C. Rules that control traffic flow to and from instances**
- D. Standards for data storage compliance**

9. What is the hard limit for the number of primary IP ranges per subnet?

- A. 1**
- B. 2**
- C. 3**
- D. 4**

10. Which of the following represents the maximum number of subnet IP ranges allowed?

- A. 1,500**
- B. 2,500**
- C. 3,100**
- D. 4,000**

Answers

SAMPLE

1. B
2. A
3. B
4. B
5. B
6. B
7. B
8. C
9. A
10. C

SAMPLE

Explanations

SAMPLE

1. Which feature distinguishes Traceroute from Ping?

- A. It measures packet loss
- B. It shows the average time for each hop**
- C. It tests bandwidth limits
- D. It performs DNS lookups

Traceroute is specifically designed to map the path that packets take from the source to the destination by sending a series of packets with incrementally increasing Time to Live (TTL) values. Each time a packet is sent out with a specific TTL, it gets decremented by each router it passes through. When the TTL reaches zero, the router discards the packet and sends back an ICMP "Time Exceeded" message. This process allows Traceroute to identify each hop along the way, recording the round-trip time for each individual hop. As a result, it displays the latency associated with each hop in the network path. In contrast, Ping primarily focuses on determining the availability of a host and measuring round-trip time between the source and a single destination address. While it can indicate packet loss if packets do not return, it does not provide detailed information about the specific hops between the source and destination. Unlike Traceroute, Ping does not break down the journey into multiple segments or hops, which is why it does not have the ability to show the average time for each hop. Overall, the characteristic of displaying the average time for each hop is what distinctly separates Traceroute from Ping.

2. If a VM has 2 or fewer vCPUs, what is the maximum number of NICs it can have?

- A. 1**
- B. 2
- C. 4
- D. 8

In Google Cloud, the limitations on virtual machines (VMs) regarding the number of network interface cards (NICs) depend on the configuration of the VM's virtual CPUs (vCPUs). For a VM with 2 or fewer vCPUs, the maximum number of NICs that it can support is limited to one. This restriction exists primarily to ensure efficient resource allocation and network throughput. When VMs are configured with fewer vCPUs, they generally handle lower workloads, which correlates with the need for only a single NIC. Utilizing additional NICs on such VMs would not significantly enhance network performance or capability, making this constraint logical in terms of resource management. When VMs are provisioned with more vCPUs, they can support more NICs. However, in this scenario, the capability is capped at one NIC for VMs with 2 or fewer vCPUs, making it essential for users to plan their VM configurations according to their performance and network needs.

3. What is the maximum response size that can be cached by Cloud CDN when byte-range requests are supported?

- A. 10 MB**
- B. 5 TB**
- C. 1 GB**
- D. 50 MB**

The maximum response size that can be cached by Cloud CDN when byte-range requests are supported is 5 TB. This allows Cloud CDN to efficiently cache and serve large objects, making it highly scalable for various applications that require the delivery of substantial amounts of data. Cloud CDN is designed to support high-throughput content delivery while ensuring that byte-range requests are efficiently handled. The 5 TB limit ensures that large media files, software distributions, or any other substantial datasets can be cached, which enhances performance and reduces latency for end-users. This capability is especially beneficial for applications dealing with high volumes of traffic and large data sizes, as it takes full advantage of Google's global distribution network to cache and serve content closer to users, making network usage more efficient and cost-effective. In contrast, the other options present lower limits that do not align with the scalable nature of Cloud CDN in handling large response sizes effectively. This significant caching capacity is crucial for optimizing content delivery on the cloud.

4. What does the term 'interconnectivity options' in a VPC refer to?

- A. Methods for enhancing data storage**
- B. Ways to connect the VPC with external networks**
- C. Protocols for internal communications**
- D. Settings for firewall configurations**

Interconnectivity options in a Virtual Private Cloud (VPC) primarily refer to the various methods available for connecting the VPC with external networks. This encompasses services and configurations that enable communication between a VPC and on-premises networks, other cloud environments, or the public internet. Such options include Cloud VPN, Interconnect, and VPC peering, which help establish secure and efficient connections. The ability to connect to other networks is crucial for hybrid cloud scenarios, where businesses utilize both cloud and on-premises resources, allowing for data exchange and seamless operations across diverse environments. Understanding these interconnectivity options is essential for designing effective network architectures in Google Cloud.

5. What does Google Cloud Armor provide to its users?

- A. Firewall protection against unauthorized access
- B. Security policies to protect applications from DDoS attacks**
- C. Encryption for data at rest
- D. User authentication services

Google Cloud Armor provides security policies designed to protect applications from Distributed Denial of Service (DDoS) attacks, making it a crucial tool for maintaining the availability and performance of web applications in the cloud. It includes features that enable users to implement robust security measures specifically tailored to mitigate the risk of DDoS threats. The service not only helps in protecting applications at the edge but also offers customizable security policies that can adapt to the specific needs of different applications. This ensures that legitimate traffic can reach your service while blocking harmful traffic aimed at disrupting service availability. In contrast, while firewall protection against unauthorized access is a critical aspect of overall cloud security, this function is more closely associated with Google Cloud's VPC Firewall rules and doesn't specifically encompass the specialized DDoS protections that Google Cloud Armor offers. Similarly, encryption for data at rest is a separate security feature that protects data stored in Google Cloud, and user authentication services are typically handled by different products, such as Identity and Access Management (IAM) tools. Therefore, the focus of Google Cloud Armor on providing security policies primarily aimed at mitigating DDoS attacks distinguishes it as the correct answer in this context.

6. Which command would you use to import a zone file into GCP?

- A. **gcloud dns import zone-file**
- B. **gcloud dns record-sets import**
- C. **gcloud dns zone import**
- D. **gcloud dns records add**

To import a zone file into Google Cloud Platform (GCP), the correct command is indeed to use `gcloud dns record-sets import`. This command is specifically designed to import DNS records from a zone file into a Google Cloud DNS managed zone. Zone files contain DNS resource record configurations, and the `record-sets import` command processes the contents of these files to create or update resource records within your specified managed zone in GCP. By utilizing this command, you efficiently handle multiple DNS records at once, streamlining the setup and management of DNS settings. The other commands provided do not serve the same function as `gcloud dns record-sets import`. They may reference operations related to DNS, but they do not perform the importation of zone files into GCP, which is critical for migrating DNS configurations or for the initial setup of a managed zone.

7. What is a limitation of the Netperf tool compared to IPerf?

- A. It cannot measure latency
- B. It has less customizable features**
- C. It does not support UDP testing
- D. It only works in Windows environments

The choice regarding the limitation of the Netperf tool relative to IPerf being that it has less customizable features is accurate. Netperf, while a powerful tool for performance measurement, has a more rigid set of options for testing compared to IPerf. IPerf is known for its flexibility and extensive configuration capabilities, allowing users to specify various parameters related to bandwidth, protocol types, buffer sizes, and more. This makes IPerf particularly well-suited for users who require fine-tuned testing scenarios and the ability to adjust settings based on specific needs. In contrast, Netperf focuses on measuring network performance with predefined benchmarks, which can limit its adaptability in certain testing environments. Users looking for extensive customization may find IPerf's capabilities better suited to their requirements. This distinction between the two tools highlights why the limitation of Netperf in terms of customizable features is a valid consideration when choosing a performance measurement tool for networking scenarios.

8. Which of the following best describes firewall rules?

- A. Protocols to manage application performance
- B. Configurations that allow secure connections to public networks
- C. Rules that control traffic flow to and from instances**
- D. Standards for data storage compliance

Firewall rules are best described as rules that control traffic flow to and from instances. These rules are essential for managing network security within Google Cloud. They specify which incoming and outgoing connections are permitted or denied, based on various criteria such as IP address ranges, protocols, and ports. This traffic control is critical for protecting cloud resources from unauthorized access while allowing legitimate traffic to flow. Firewall rules enable network engineers to enforce security policies at the network level, ensuring that only designated traffic is allowed through the firewall, which is essential for maintaining a secure infrastructure in cloud environments. This functionality is vital for both protecting sensitive data and ensuring that applications operate correctly by allowing necessary communications while restricting unwanted access.

9. What is the hard limit for the number of primary IP ranges per subnet?

- A. 1**
- B. 2**
- C. 3**
- D. 4**

The hard limit for the number of primary IP ranges per subnet in Google Cloud is indeed one. Each subnet in a Virtual Private Cloud (VPC) can only have a single primary IP range allocated to it. This restriction simplifies the management of IP addressing within subnets and ensures that there is a clear and unambiguous allocation of IP addresses within the network. When a subnet is created, a primary IP range is designated, which defines the IP address space from which instances within that subnet can be assigned addresses. The decision to limit this to one primary range is primarily to avoid complications associated with overlapping IP ranges, which could lead to routing issues and conflicts in network operations. Secondary IP ranges can be added to a subnet, allowing for more complex scenarios such as aliasing IP addresses for services like Kubernetes clusters, but they are not considered primary IP ranges. This distinction highlights the simplicity of the primary IP range concept, which is essential for maintaining clear network structure and integrity in Google Cloud environments.

10. Which of the following represents the maximum number of subnet IP ranges allowed?

- A. 1,500**
- B. 2,500**
- C. 3,100**
- D. 4,000**

In Google Cloud, the maximum number of IP ranges that can be assigned to a single subnet within a Virtual Private Cloud (VPC) network is 3,100. This limit is important because it allows for efficient network design and IP address management within the environment. When architecting VPC networks to support applications and services, understanding this limitation is crucial to ensure that the network can efficiently scale and accommodate the necessary amount of resources while avoiding potential issues concerning IP address exhaustion or misconfiguration. The significance of the maximum number of subnet IP ranges is also associated with how resources are deployed and how network traffic is routed within the Google Cloud environment. A higher number of available subnet ranges facilitates more granular control over network segmentation, allowing for better organization of resources based on security or operational requirements. Understanding this capacity helps network engineers in planning and segmenting their Google Cloud networks appropriately, ensuring that they take full advantage of the available resources while adhering to the constraints imposed by the platform.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

<https://googlecloudprofessionalcloudnetengr.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE